

WHAT IS CLAIMED IS:

1. A system, comprising:
 - 5 a node including a memory, an active device, an interface to an inter-node network, and an address network coupling the memory, the active device, and the interface;

an additional node coupled to the node by the inter-node network;
 - 10 wherein the active device is configured to initiate a transaction to gain an access right to a coherency unit by sending an address packet on the address network;
 - 15 wherein if the transaction cannot be satisfied within the node, the interface is configured to send a coherency message to the additional node via the inter-node network;
 - 20 wherein in response to receiving an additional coherency message from the additional node via the inter-node network, the interface is configured to send data corresponding to the coherency unit to the active device, wherein the active device gains the access right upon receipt of the data.
2. The system of claim 1, wherein the address packet is a read-to-own packet and the
25 access right is a write access right.
3. The system of claim 2, wherein an additional interface included in the additional node is configured to receive the coherency message and to responsively send a proxy-read-to-own modified packet on an additional address network included in the additional
30 node.

4. The system of claim 3, wherein an additional active device included in the additional node is configured to lose an ownership responsibility for the coherency unit upon receipt of the proxy-read-to-own modified packet.

5

5. The system of claim 4, wherein the additional active device in the additional node is configured to send a data packet corresponding to the coherency unit on an additional data network included in the additional node in response to receiving the proxy read-to-own modified packet, wherein the additional active device transitions an access right to the coherency unit upon sending the data packet.

10

6. The system of claim 5, wherein the additional interface is configured to send the additional coherency message on the inter-node network in response to receiving the data packet sent by the additional active device, wherein the additional coherency message includes a copy of the coherency unit.

15

7. The system of claim 2, wherein in response to receiving the coherency message, the additional interface is configured to send a proxy read-to-own packet on an additional address network included in the additional node;

20

wherein in response to receiving the proxy read-to-own packet, a plurality of additional active devices included in the additional node are configured to invalidate an access right to the coherency unit;

25

wherein the additional interface is further configured to send an invalidating coherency message on the inter-node network to one or more other nodes that include active devices having a read access right to the coherency unit.

8. The system of claim 7, wherein in response to the invalidating coherency message, interfaces included in each of the one or more other nodes are configured to

30

send a proxy-invalidate coherency message on a respective address network and to send an acknowledgment coherency message to the node in response to receiving the proxy-invalidate coherency message on the respective address network.

5 9. The system of claim 8, wherein the interface included in the node is configured to not provide the data corresponding to the active device until an acknowledgment coherency message has been received from each of the one or more other nodes.

10 10. The system of claim 8, wherein an additional memory included in the additional node is configured to send data corresponding to the coherency unit to the additional interface in response to receiving the proxy read-to-own packet; and

wherein the additional interface is configured to include the data in the additional coherency message sent to the interface included in the node.

15

11. The system of claim 1, wherein the additional node is a home node for the coherency unit.

20 12. The system of claim 1, wherein the interface is configured to send the coherency message in response to an indication that the transaction cannot be completed within the node, wherein the indication is an address of the coherency unit, wherein the address of the coherency unit does not map to any memory included in the node.

25 13. The system of claim 1, wherein the interface is configured to send the coherency message in response to an indication that the transaction cannot be completed within the node, wherein the indication is a report corresponding to the address packet generated by the memory in response to a global access state of the coherency unit in the node.

14. The system of claim 1, wherein the active device does not implement different operating modes based on whether a given coherency unit maps or does not map to a memory within the node.

5 15. The system of claim 1, wherein the active device does not implement different operating modes based on whether the active device is included in a single node system or a multi-node system.

10 16. The system of claim 1, wherein the active device is configured to send a same type of address packet containing a same command encoding to initiate the transaction for the coherency unit in a single node system as the active device is configured to send to initiate the transaction for the coherency unit in a multi-node system.

15 17. The system of claim 1, wherein the active device is configured to send a same type of address packet containing a same command encoding to initiate the transaction for the coherency unit if the coherency unit does not map to any memory included in the node as the active device is configured to send to initiate the transaction for the coherency unit if the coherency unit does map to the memory included in the node.

20 18. The system of claim 1, wherein the active device is configured to initiate the transaction using a same type of address packet regardless of whether the transaction can be satisfied within the node.

25 19. A system, comprising:

a node including a memory subsystem, an active device, an interface to an inter-node network, and an address network coupling the memory subsystem, the active device, and the interface;

30 an additional node coupled to the node by the inter-node network;

5 wherein the active device is configured to initiate a transaction to gain an access right to a coherency unit by sending an address packet including a command encoding on the address network, wherein the coherency unit maps to the memory subsystem;

10 wherein the active device is configured to initiate an additional transaction to gain the access right to a different coherency unit by sending an additional address packet containing the command encoding on the address network, wherein the different coherency unit does not map to any memory subsystem in the node;

15 wherein the interface is configured to send a coherency message corresponding to the additional address packet to the additional node via the inter-node network in response to receiving the additional address packet.

20. The system of claim 19, wherein the command encoding is a read-to-own command encoding, wherein the access right is a write access right.

20 21. The system of claim 19, wherein the command encoding is a read-to-share command encoding, wherein the access right is a read access right.

22. A system, comprising:

25 a node including a memory subsystem, an active device, an interface to an inter-node network, and an address network coupling the memory subsystem, the active device, and the interface;

an additional node coupled to the node by the inter-node network;

30

wherein the active device is configured to initiate a coherency transaction to gain an access right to a coherency unit by sending a type of address packet on the address network, wherein the coherency unit maps to the memory subsystem;

5

wherein if the coherency transaction cannot be satisfied within the node, the interface is configured to send a coherency message to the additional node via the inter-node network;

10

wherein the active device does not initiate coherency transactions dependent on an indication as to whether the active device is currently included in a single-node computer system or a multi-node computer system.

23. A method for use in a multi-node system comprising a node coupled to an additional node by an inter-node network, the method comprising:

15

an active device included in the node initiating a transaction to gain an access right to a coherency unit by sending an address packet on an address network included in the node;

20

an interface to the inter-node network receiving the address packet via the address network;

25

if the transaction cannot be completed within the node, the interface sending a coherency message corresponding to the address packet to the additional node via the inter-node network;

30

in response to receiving an additional coherency message from the additional node via the inter-node network, the interface sending data corresponding to the coherency unit to the active device;

the active device gaining the access right upon receipt of the data.

24. The method of claim 23, wherein the address packet is a read-to-own packet and
5 the access right is a write access right.

25. The method of claim 24, further comprising an additional interface to the inter-
node network included in the additional node receiving the coherency message and
responsively sending a proxy-read-to-own modified packet on an additional address
10 network included in the additional node.

26. The method of claim 25, further comprising an additional active device included
in the additional node losing an ownership responsibility for the coherency unit upon
receipt of the proxy-read-to-own modified packet.

15

27. The method of claim 26, further comprising:

the additional active device in the additional node sending a data packet
corresponding to the coherency unit on an additional data network
20 included in the additional node in response to receiving the proxy read-to-
own modified packet; and

the additional active device transitioning an access right to the coherency unit to
invalid upon sending the data packet.

25

28. The method of claim 27, further comprising the additional interface sending the
additional coherency message on the inter-node network in response to receiving the data
packet sent by the additional active device, wherein the additional coherency message
includes a copy of the coherency unit.

30

29. The method of claim 24, further comprising:

in response to receiving the coherency message, the additional interface sending a proxy read-to-own packet on an additional address network included in the additional node;

in response to receiving the proxy read-to-own packet, a plurality of additional active devices included in the additional node invalidating an access right to the coherency unit;

the additional interface sending an invalidating coherency message on the inter-node network to one or more other nodes that include active devices having a read access right to the coherency unit.

30. The method of claim 29, further comprising:

in response to the invalidating coherency message, interfaces included in each of the one or more other nodes sending a proxy-invalidate coherency message on a respective address network and sending an acknowledgment coherency message to the node in response to receiving the proxy-invalidate coherency message on the respective address network.

31. The method of claim 30, further comprising the interface included in the node not providing the data corresponding to the active device until an acknowledgment coherency message has been received from each of the one or more other nodes.

32. The method of claim 30, further comprising:

an additional memory included in the additional node sending data corresponding to the coherency unit to the additional interface in response to receiving the proxy read-to-own packet; and

5 the additional interface including the data in the additional coherency message sent to the interface included in the node.

33. The method of claim 23, wherein the additional node is a home node for the coherency unit.

10

34. The method of claim 23, wherein said sending the coherency message comprises the interface sending the coherency message in response to an indication that the transaction cannot be completed within the node, wherein the indication is an address of the coherency unit, wherein the address of the coherency unit does not map to any

15 memory included in the node.

35. The method of claim 23, wherein said sending the coherency message comprises the interface sending the coherency message in response to an indication that the transaction cannot be completed within the node, wherein the indication is a report corresponding to the address packet generated by the memory in response to a global access state of the coherency unit in the node.

20

36. The method of claim 23, wherein the active device does not implement different operating modes based on whether a given coherency unit maps or does not map to a memory within the node.

25

37. The method of claim 23, wherein the active device does not implement different operating modes based on whether the active device is included in a single node system or a multi-node system.

30

38. The method of claim 23, wherein said initiating the transaction comprises the active device sending a same type of address packet containing a same command encoding to initiate the transaction for the coherency unit in a single node system as the active device is configured to send to initiate the transaction for the coherency unit in a multi-node system.

39. The method of claim 23, wherein said initiating the transaction comprises the active device sending a same type of address packet containing a same command encoding to initiate the transaction for the coherency unit if the coherency unit does not map to any memory included in the node as the active device is configured to send to initiate the transaction for the coherency unit if the coherency unit does map to the memory included in the node.

40. The method of claim 23, wherein said initiating the transaction comprises the active device sending a same type of address packet regardless of whether the transaction can be satisfied within the node.